Cognitive Robotics

Giulio Sandini
Robotics, Brain and Cognitive Sciences Unit
Istituto Italiano di Tecnologia
Via Melen 83
16152 - Genova
Italy
giulio.sandini@iit.it

Alessandra Sciutti
COgNiTive Architecture for Collaborative Technologies Unit
Istituto Italiano di Tecnologia
Via Melen 83
16152 - Genova
Italy
alessandra.sciutti@iit.it

David Vernon
Institute for Artificial Intelligence
University of Bremen,
Am Fallturm 1
28359 Bremen
Germany
dvernon@uni-bremen.de

Synonyms

Cognitive embodied systems, artificial cognitive systems

Definition

The word cognition derives from the Latin verb *cognosco*, a composition of *con* (meaning related to) and *gnosco* (to know). Cognitive robotics, then, is the branch of robotics where knowledge plays a central role in supporting action selection, execution, and understanding. It focuses on designing and building robots that have the ability to learn from experience and from others, commit relevant knowledge and skills to memory, retrieve them as the context requires, and flexibly use this knowledge to select appropriate actions in the pursuit of their goals, while anticipating the outcome of those actions when doing so. Cognitive robots can use their knowledge to reason about their actions and the actions of those with whom they are interacting, and thereby modify their behavior to improve their overall long-term effectiveness. In short, cognitive robots are capable of flexible, context-sensitive action, knowing what they are doing and why they are doing it.

Overview

There are two aspects to the goal of building a cognitive robot (Krichmar 2012). One is to gain a better understanding of cognition in general — the so-called synthetic methodology — and the other is to build systems that have capabilities that are rarely found in technical artifacts (i.e. artificial systems) but are commonly found in humans and some animals. The motivation for the first is a principled one, the motivation for the second is a practical one (Vernon 2017).

In fact, the field of cognitive robotics combines insights and methods from robotics, artificial intelligence, and the cognitive and biological sciences. It emphasizes bio-inspired human-like and animal-like behavior and intelligence, system-level integration of a range of cognitive abilities, including sensorimotor skills, knowledge representation and reasoning, and social interaction. It takes an interdisciplinary approach, including cognitive science, cognitive neuroscience, cognitive psychology, and biology (Cangelosi and Asada, in press).

Key Research Findings

The Role of Cognition

Cognition enables a robot to pursue goals in a flexible, experience-based and context-sensitive manner. It allows robots to work autonomously in everyday environments, anticipating outcomes when selecting the actions they will perform and adapting to changes and unforeseen situations. In particular, cognitive robots can recruit the motions that are most likely to result in the successful achievement of each action, predictively control these motions and adapting them if they fail. It also means they can reason about their actions and the actions of those with whom they are interacting, and thereby modify their behavior to improve their overall long-term effectiveness in manipulation, navigation, and social interaction. A cognitive robot anticipates the outcome of its own goal-directed actions and also the actions executed by others, and it understands their goals and intentions.

Action itself is guided by prospection, i.e. the future anticipated state of the world, both when selecting the action and when carrying it out. A cognitive robot can also adapt to changing circumstances, adjusting existing action policies and creating new action policies when required. It does all this by dynamically recruiting several core cognitive abilities: perception, action, learning, adaptation, anticipation, motivation, autonomy, internal simulation, attention, action selection, memory, reasoning, and meta-reasoning (Vernon 2014; Vernon, von Hofsten, and Fadiga 2016; Kotseruba and Tsotsos 2020).

The list is not exhaustive and one could add other functionalities, as for instance development (Asada et al. 2009; Cangelosi and Schlesinger 2015; Vernon, von Hofsten, and Fadiga 2016). In fact, it has been proposed that for robots to learn and think like humans they will need developmental processes which build from a basis of core knowledge of number, space, physics, and psychology (or agency), and that can generalize knowledge by exploiting capacities in the architecture for compositionality and learning-to-learn. Also, cognitive robots will need a capacity to build causal models which facilitate abductive inference (i.e. explanation) and counterfactual reasoning, as well as prospection.

Cognition is essential for effective interaction with humans because it enables a robot to infer the goals and intentions of the person with whom it is interacting and thereby allows it to behave in a helpful manner, cooperating to assist the humans achieve their goals or collaborating to achieve joint goals. Furthermore, humans have an innate propensity to engage with other cognitive agents. so if a robot has a capacity for cognition it fosters human robot interaction and leads to bi-directional engagement, one of the key elements of social robotics.

An important feature of cognitive robotics is the focus on prospection to augment immediate sensorimotor experience, both when navigating and manipulating objects in the robot's environment and when interacting with humans. Cognitive robots are able to carry out tasks effectively by anticipating the effects of their own actions as well as the actions of the humans around them on the basis of shared experience. Being able to view the world from another person's perspective, a cognitive robot can anticipate that person's intended actions and needs and, consequently, it can interact safely while performing tasks in everyday situations. This applies both during direct interaction (e.g. a robot assisting a customer in a supermarket) and during indirect interaction (e.g. a robot stacking shelves while customers are moving around it while shopping). In this respect "anticipation" means being able to infer the goal of other's actions on the basis of shared knowledge built through past experience. In short, cognitive robots know what they are doing (the expected outcome of their actions) and understand what others are doing (the expected outcome of others' actions).

Core Cognitive Abilities

Cognitive robots achieve their goals by perceiving their environment, paying attention to the events that matter, anticipating the need for some action, planning what to do, anticipating the outcome as they execute the action, learning from the resultant interaction, and adapting to change.

Perception makes use of many sensory modalities, e.g. vision, audition, and haptic (tactile and kinaesthetic). Attention operates in different ways, selecting priority features or objects, restricting what to pay attention to and where to pay attention to, and suppressing features, objects, or locations that are deemed to be not relevant (Tsotsos 2011). It is worth stressing the fact that the effectiveness of attentional shift is greatly improved by anticipation, by devoting more sensory and processing power where "something is going to happen" instead of relying only on real-time sensory feedback. Anticipation, also referred to as prospection, is often associated with achieving a goal. There are four modes of operation: simulation, prediction, intention, and planning (Szpunar et al. 2014). Planning is sometimes effected by reasoning about the current state of the world or anticipated futures states. It exploits memories of past experience, held in episodic memory, and knowledge of the world, held in semantic memory. Anticipating the outcome of a possible action can refer to the actions of the robot itself or the actions of other agents, human or robot. Learning from actions means that future actions can be more effective or more efficient. This is often based on reasoning. It is sometimes referred to as metacognition or meta-reasoning when the focus is on improving the cognitive or reasoning process. Adaptation is also achieved through learning.
In this case, the result of learning is a new action policy rather than the improvement of an existing action policy.

In cognitive robotics, perception and action both become inherently cognitive, in that they acquire meaning only when, acting together, they serve as an enabler for construction and validation of predictions, which in turn guide perception and action selection, through planning and internal simulation. This “exploration-prediction” cycle, which expands the action-perception one, involves building, or learning, models of the world and the person with whom the robot is interacting, and adapting these models when necessary to improve their predictive power.

Core Cognitive Abilities are Coordinated in a Cognitive Architecture

The chief characteristic of a cognitive agent is the ability to act effectively in a world that is uncertain, under-specified, dynamic, possibly cooperating with other cognitive agents. To achieve goals adaptively and robustly in these circumstances requires a complex system that can construct models of the way the world works, use them to guide actions prospectively, and update them dynamically as the system continuously expands its knowledge and experience through its interactions. A cognitive architecture is the way we specify what is required to achieve this.

A cognitive architecture is a software framework that integrates all the elements required for a system to exhibit the characteristic attributes of a cognitive agent. It determines the overall structure and organization of a cognitive system, including the component parts or modules (Sun 2004), the dynamic relations between these modules, and the essential algorithmic and representational details within them (Langley 2006, Langley and Choi, 2006). The architecture specifies the formalisms for knowledge representation and the types of memories used to store this knowledge, the processes that act upon it, and the learning mechanisms that acquire and updates it.

A cognitive architecture integrates the core cognitive abilities so that these abilities can be dynamically coordinated, thereby allowing the agent to exhibit flexible context-sensitive behavior, prospectively selecting and controlling the actions that are required to achieve given goals. A cognitive architecture should also be able to develop autonomously so that its performance improves over time with experience.

The majority of cognitive architectures focus on modelling human cognition. Cognitive architectures specifically designed for cognitive robots include ISAC (Kawamura et al. 2009), iCub (Sandini et al. 2007; Vernon et al. 2007), CRAM (Beetz et al. 2010; Mösenlechner 2016), and a self-verifying developmental cognitive architecture (Wieser and Cheng, 2018). For a more detailed discussion and a complete list of references, please refer to the entry *Cognitive Architecture* in this volume.

Example Applications

Natural Human Robot Interaction

The iCub humanoid robot is used as a platform implemented for open research in cognitive robotics (Sandini, Metta, and Vernon, 2007; Vernon, Matta and Sandini, 2007, Metta et al. 2010; www.icub.org). The iCub is especially suitable for applications that require natural human robot interaction and target collaboration between a human and a cognitive robot. Fig.1 shows a sequence of pictures that depicts a situation in which the iCub is interacting with a human, reading her intention to get her phone from her bag, and alerting her to the fact that it is on the desk, hidden from her by the laptop. Note that this sequence has been staged to illustrate the desired capabilities of a cognitive robot and has not yet been implemented.

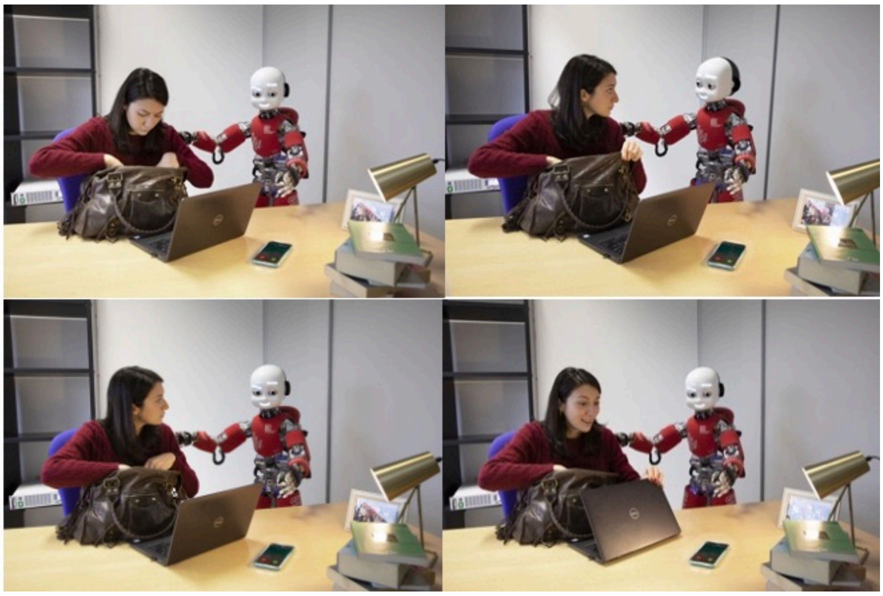

Fig. 1. The iCub humanoid robot: a platform for open research in cognitive robotics.

Robot Manipulation in Everyday Activities

One of the goals of cognitive robotics is for humans to be able to give a robot some task to do by stating that task in the same terms they would use if they were talking to another human being, conveying just the essence of the goal without have to specify exactly how that task is the be carried out.

Fig. 2. A PR2 robot demonstrating cognition-enabled manipulation by using knowledge and reasoning to determine the motions required to pour popcorn from a saucepan.

For example, one might say to a cognitive robot, “put the dishes in the dishwasher”. The cognitive robot then has to map this implicit task description onto a series of explicit actions and related manipulation motions so that the task is accomplished successfully. This capability is what the CRAM cognitive architecture (Beetz et al. 2010) makes possible.[1] Fig 2. shows a PR2 robot in the process of pouring popcorn from a saucepan during a demonstration of cognitively-enabled robot manipulation using CRAM.

Future Directions for Research

The complexity of cognition and the scale of the challenge of attempting to generate a robot that is capable of cognitive interactive behavior be it by data-driven deep learning or classical computational modelling or some mix of the two, is apparent.

Still today we have not a clear answer to the question: "[…], accepting that cognition is achieved in large part through development and learning, what operational architecture is needed to get the cognitive development started?" (Auer et al., 2005). What is the minimal necessary and sufficient phylogenetic configuration a cognitive system for it (a) to exhibit cognitive skills and (b) to develop these cognitive skills to a level that enables interaction with humans?

There are many factors – and many research challenges – to consider when answering this question. There is a need for more realistic perceptual capabilities that can operate in adverse conditions with noise and uncertainty, using context to improve performance. More specifically, cognitive systems, in general, and perception systems, in particular, need to be able to generalize more effectively than they can at present, hypothesizing models that extend beyond the distribution in which have been trained.

The cognitive architectures that form the software platform of cognitive robots need to facilitate more natural communication with humans, allowing them to infer a human's intentions and emotional states, to engage in adaptive, personalized interaction, to read body language, e.g. gestures and facial expressions, to engage in natural turn-taking, and facilitate human-robot joint action. Example cognitive architectures that focus on these aspects of cognitive human-robot interaction include (Lemaignan et al. 2017; Sandini et al. 2018; Tanevska et al. 2019).

Computational models of episodic memory have not received significant attention, especially for life-long learning, despite the fact that its existence and importance has been widely recognized (Kotseruba and Tsotsos 2020). Notable exceptions include the CRAM cognitive architecture (Beetz, Mösenlechner, and Tenorth 2010; Moösenlechner 2016) and the iCub neural framework for episodic memory (Mohan, Sandini, and Morasso 2014).

Deep learning (Schmidhuber 2014; Goodfellow, Bengio, and Courville 2016) has not yet made a significant impact on cognitive architectures (Kotseruba and Tsotsos 2020). This will almost certainly change, giving rise to new architectural requirements, e.g. deep developmental

[1] The achievement of CRAM's ability to map underdetermined instructions to the explicit robot motions required to carry out the actions successfully was developed over the next ten years, most recently at the EASE (Everyday Activity Science and Engineering) interdisciplinary research center (www.ease-crc.org), University of Bremen, Germany.

robotics architectures (Sigaud and Droniou 2016) and a reconciliation of deep learning with symbolic artificial intelligence (Garnelo and Shanahan 2019).

Additionally, if it is true that social interaction is the default mode in which the brain operates (Hari, Henriksson, Malinen, & Parkkonen, 2015) a new approach to the modeling of the mind is needed where social interaction is not seen as the mere superposition of two individual cognitive agents (Vernon, 2014, Sandini et al. 2018) but is actively exploited to build a shared experience used to anticipate each other intentions as illustrated with an example in Figure 1. Hence, inherently social architecture and social cognitive models need to be developed for cognitive robotics. This calls for help from disciplines that go beyond the ones traditionally involved in the design of cognitive architecture: neuroscience, AI, robotics and cognitive science. Indeed, sociologist, anthropologists, philosophers, economists and artists, among others, have a grasp of the social milieu that could help shed a novel light on the current issue.

Cross-References

Cognitive Architecture
Developmental Robotics
iCub
Human-Centered Robotics
Social Robotics
Soft Robotics
Cognitive Human-Robot Interaction

References

Asada, M., K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M.i Ogino, and C. Yoshida. Cognitive developmental robotics: a survey. IEEE Transactions on Autonomous Mental Development, Vol.1, No.1, pp.12--34, 2009

Auer, P., A. Billard, H. Bischof, I. Blochand, P. Boettcherand, H. B. H. Buxton, H. Christensen, T. Cohn, P. Courtney, A. Crookell, J. Crowley, S. Dickinson, C. Eberst, J. Olof Eklundh, B. Fisher, W. F˙orstner, J. Gilby, G. Granlund, V. Hlavac, J. Kittler, W. Kropatsch, A. Leonardis, J. Little, G. Metta, H.-H. Nagel, B. Nebel, B. Neumann, H. Niemann, L. Paletta, A. Pinz, F. Pirri, G. Sagerer, G. Sandini, B. Schiele, R. Simpson, G. Sommer, J. Tsotsos, M. Thonnat, D. Vernon, and M. Vincze (2005). A Research Roadmap of Cognitive Vision. ECVision: European Network for Research in Cognitive Vision Systems. IST Project IST-2001-35454. https://doi.org/10.1.1.154.721.

Beetz, M., L. Mösenlechner, and M. Tenorth (2010, October). CRAM – A Cognitive Robot Abstract Machine for Everyday Manipulation in Human Environments. In IEEE/RSJ International Conference on Intelligent Robots and Systems, Taipei, Taiwan, pp. 1012–1017.

Cangelosi, A. and Asada, M. Cognitive Robotics, Chapter 1, MIT Press, in press.

Hari, R., Henriksson, L., Malinen, S., & Parkkonen, L. (2015). Centrality of Social Interaction in Human Brain Function. Neuron, 88(1), 181–193. https://doi.org/10.1016/j.neuron.2015.09.022

Cangelosi, A. and M. Schlesinger (2015). Developmental Robotics: From Babies to Robots. Cambridge MA: MIT Press.

Garnelo, M. and M. Shanahan (2019). Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. Current Opinion in Behavioral Sciences 29, 17–23.

Goodfellow, I., Y. Bengio, and A. Courville (2016). Deep Learning. MIT Press.

Kawamura, K., S. M. Gordon, P. Ratanaswasd, E. Erdemir, and J. F. Hall (2008). Implementation of cognitive control for a humanoid robot. International Journal of Humanoid Robotics 5 (4), 547–586.

Kieras, D. and D. Meyer (1997). An overview of the epic architecture for cognition and performance with application to human-computer interaction. Human-Computer Interaction 12(4).

Kotseruba, I. and J. Tsotsos (2020). 40 years of cognitive architectures: core cognitive abilities and practical applications. Artificial Intelligence Review 53(1), 17 – 94.

Krichmar, J. L. (2012). Design principles for biologically inspired cognitive architectures. Biologically Inspired Cognitive Architectures 1, 73–81.

Langley, P. (2006). Cognitive architectures and general intelligent systems. AI Magazine 27(2), 33–44.

Langley, P. and D. Choi (2006). A unified cognitive architecture for physical agents. In Proceedings of the Twenty-First AAAI Conference on Artificial Intelligence, Boston. AAAI Press.

Lemaignan, S., M. Warnier, E. Akin Sisbot, A. Clodic, and R. Alami (2017). Artificial cognition for social human-robot interaction: an implementation. Artificial Intelligence 247, 45–69.

Metta, G., L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor, A. Bernardino, and L. Montesano, The iCub humanoid robot: an open-systems platform for research in cognitive development. Neural Netw, 2010. 23(8-9): p. 1125-34.

Mohan, V., G. Sandini, and P. Morasso (2014). A neural framework for organization and flexible utilization of episodic memory in cumulatively learning baby humanoids. Neural Computation 26, 1–43.

Mösenlechner, L. (2016). The Cognitive Robot Abstract Machine: A Framework for Cognitive Robotics. Ph. D. thesis, Technical University of Munich.

Sandini, G., G. Metta, and D. Vernon (2007). The iCub cognitive humanoid robot: An open-system research platform for enactive cognition. In M. Lungarella, F. Iida, J. C. Bongard, and R. Pfeifer (Eds.), 50 Years of AI, Volume LNAI 4850, pp. 359–370. Heidelberg: Springer.

Sandini, G., V. Mohan, A. Sciutti, and P. Morasso (2018). Social cognition for human-robot symbiosis - challenges and building blocks. Frontiers in Neurorobotics 12, 1–19.

Schmidhuber, J. (2014). Deep learning in neural networks: An overview. arXiv preprint (arXiv:1404.7828 v2).

Sigaud, O. and A. Droniou (2016). Towards deep developmental learning. IEEE Transaction on Cognitive and Developmental Systems 8(2), 99– 114.

Sun, R. (2004). Desiderata for cognitive architectures. Philosophical Psychology 17(3), 341–373.

the Clarion Cognitive Architecture. Oxford University Press.

Szpunar, K. K., R. N. Spreng, and D. L. Schacter (2014). A taxonomy of prospection: introducing an organizational framework for future-oriented cognition. PNAS 111(52), 18414–18421.

Tanevska, A., F. Rea, G. Sandini, L. Canamero, and A. Sciutti (2019). A cognitive architecture for socially adaptable robots. In Proceedings of the Joint IEEE 9th International

Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob), pp. 195–200.
Tsotsos, J. K. (2011). A Computational Perspective on Visual Attention. Cambridge MA: MIT Press.
Vernon, D. (2010). Enaction as a conceptual framework for development in cognitive robotics. Paladyn Journal of Behavioral Robotics 1(2), 89–98.
Vernon, D. (2014). Artificial Cognitive Systems — A Primer. Cambridge, MA: MIT Press.
Vernon, D. (2017). Two ways (not) to design a cognitive architecture. In R. Chrisley, V. C. Müller, Y. Sandamirskaya, and M. Vincze (Eds.), Proceedings of EUCognition 2016, Cognitive Robot Architectures, Volume CEUR-WS Vol-1855, Vienna, pp. 42–43. European Society for Cognitive Systems.
Vernon, D., G. Metta, and G. Sandini (2007). The iCub cognitive architecture: Interactive development in a humanoid robot. In Proceedings of IEEE International Conference on Development and Learning (ICDL), Imperial College, London.
Vernon, D., C. von Hofsten, and L. Fadiga (2016). Desiderata for developmental cognitive architectures. Biologically Inspired Cognitive Architectures 18, 116–127.
Wieser, E. and G. Cheng (2018). A self-verifying cognitive architecture for robust bootstrapping of sensory-motor skills via multipurpose predictors. IEEE Transaction on Cognitive and Developmental Systems 10(4), 1081–1095. Adams S, Arel I, Bach J, Coop R, Furlan R, Goertzel B, Hall JS, Samsonovich A, Scheutz M, Schlesinger M, Shapiro SC. (2012). Mapping the landscape of human-level artificial general intelligence. AI magazine 33(1):25-42.